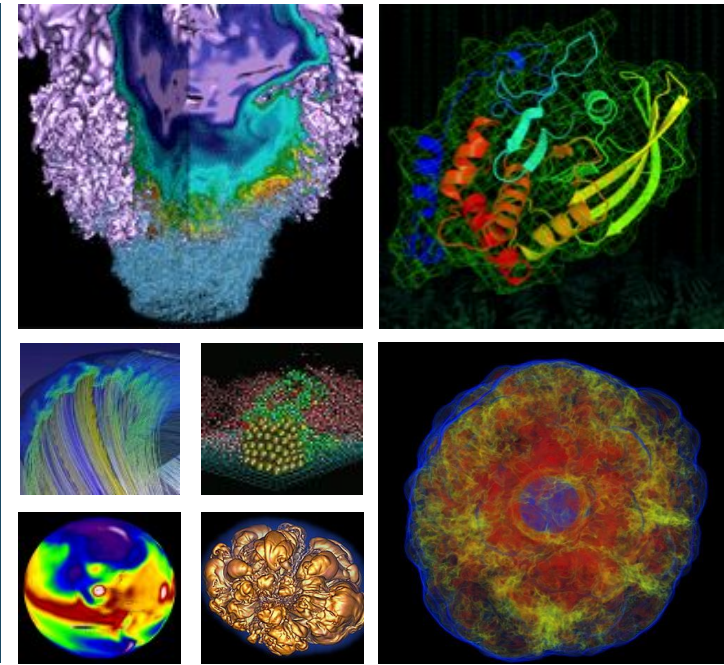


Cori Application Readiness Strategy - NESAP



Jack Deslippe
Acting Lead NERSC Apps
Performance Group
April, 2016

Edison (Ivy-Bridge):

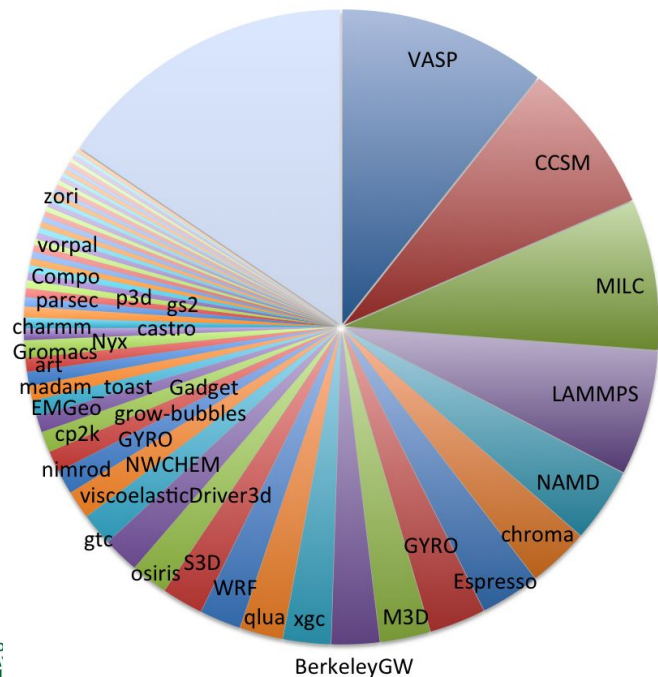
- 12 Cores Per CPU
- 24 Virtual Cores Per CPU
- 2.4-3.2 GHz
- Can do 4 Double Precision Operations per Cycle (+ multiply/add)
- 2.5 GB of Memory Per Core
- ~100 GB/s Memory Bandwidth

Cori (Knights-Landing):

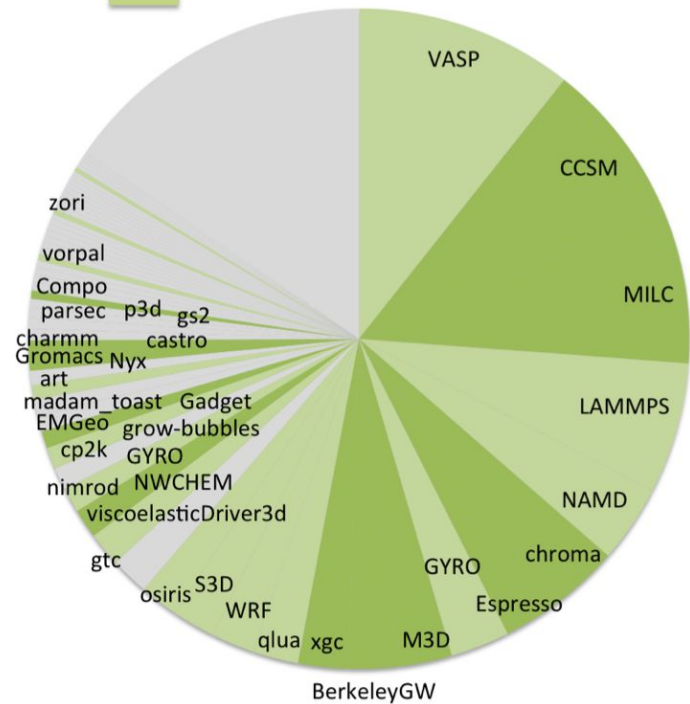
- Up to 72 Physical Cores Per CPU
- Up to 288 Virtual Cores Per CPU
- Lower GHz
- Can do 8 Double Precision Operations per Cycle (+ multiply/add)
- < 0.3 GB of HBM Memory Per Core
< 2 GB of DDR Memory Per Core
- Fast memory has ~ 5x DDR4 bandwidth

Code Coverage

**Breakdown of Application Hours
on Hopper and Edison**



NESAP Tier-1, 2 Code
 NESAP Proxy Code or Tier-3 Code



Resources for Code Teams



- **Early access to hardware**
 - Access to Babbage (KNC cluster) and early “white box” test systems expected in 2015
 - Early access and significant time on the full Cori system
- **Technical deep dives**
 - Access to Cray and Intel staff on-site staff for application optimization and performance analysis
 - Multi-day deep dive (‘dungeon’ session) with Intel staff at Oregon Campus to examine specific optimization issues
- **User Training Sessions**
 - From NERSC, Cray and Intel staff on OpenMP, vectorization, application profiling
 - Knights Landing architectural briefings from Intel
- **NERSC Staff as Code Team Liaisons (Hands on assistance)**
- **8 Postdocs**

NESAP Postdocs



Taylor Barnes
Quantum ESPRESSO



Brian Friesen
Boxlib



Andrey Ovsyannikov
Chombo-Crunch



Mathieu Lobet
WARP

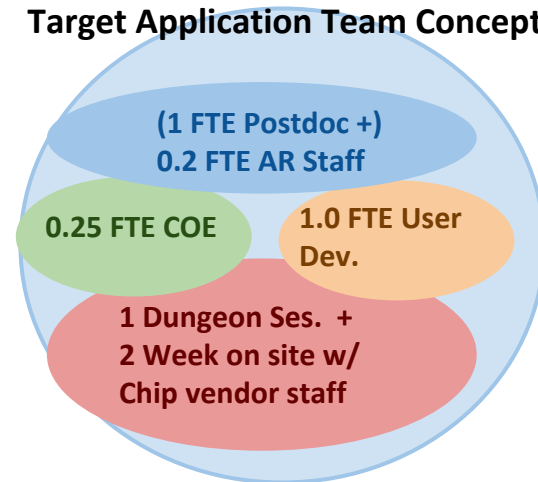


Tuomas Koskela
XGC1

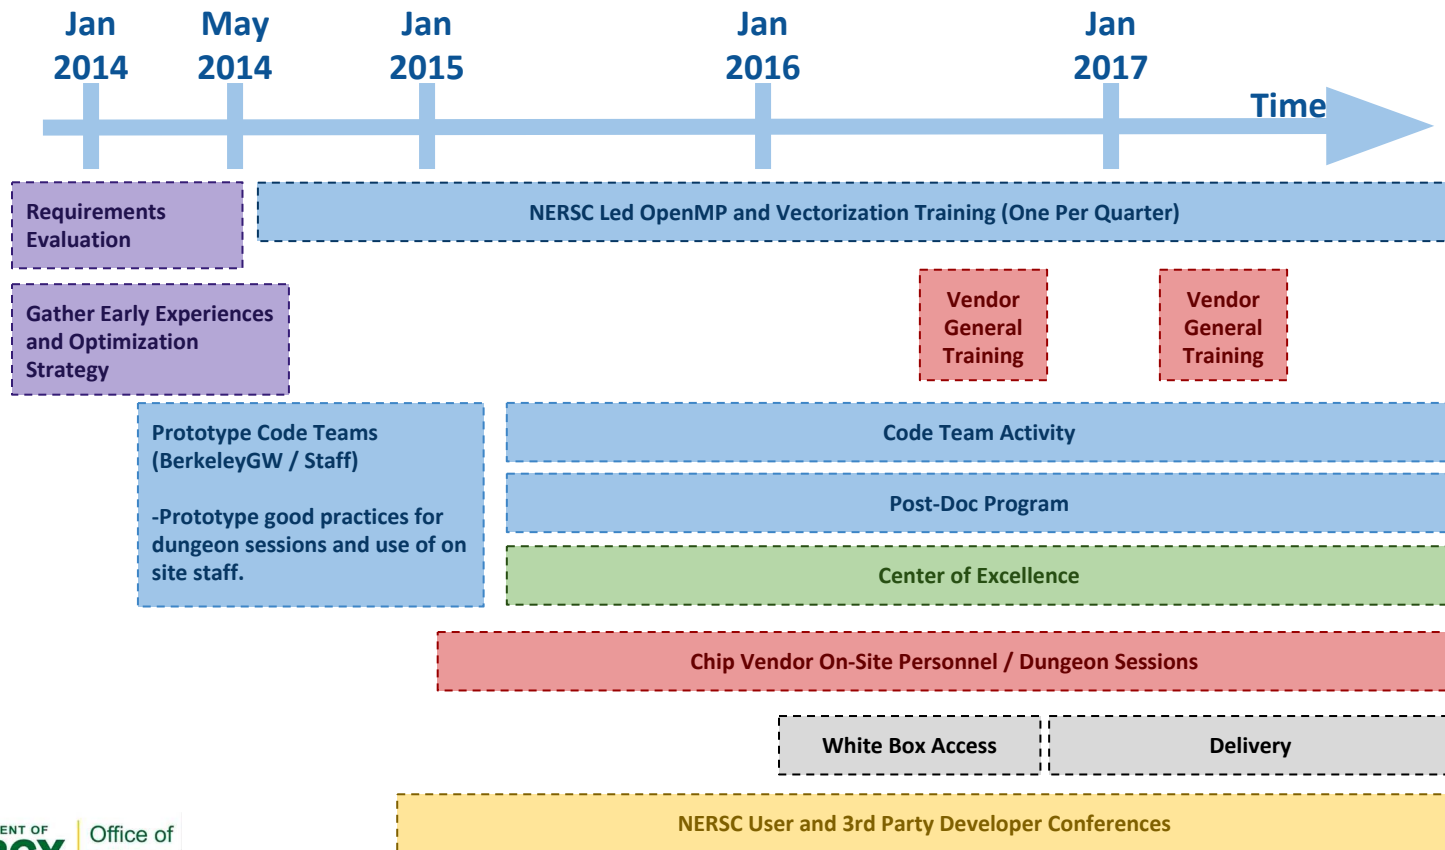


Tareq Malas
EMGeo

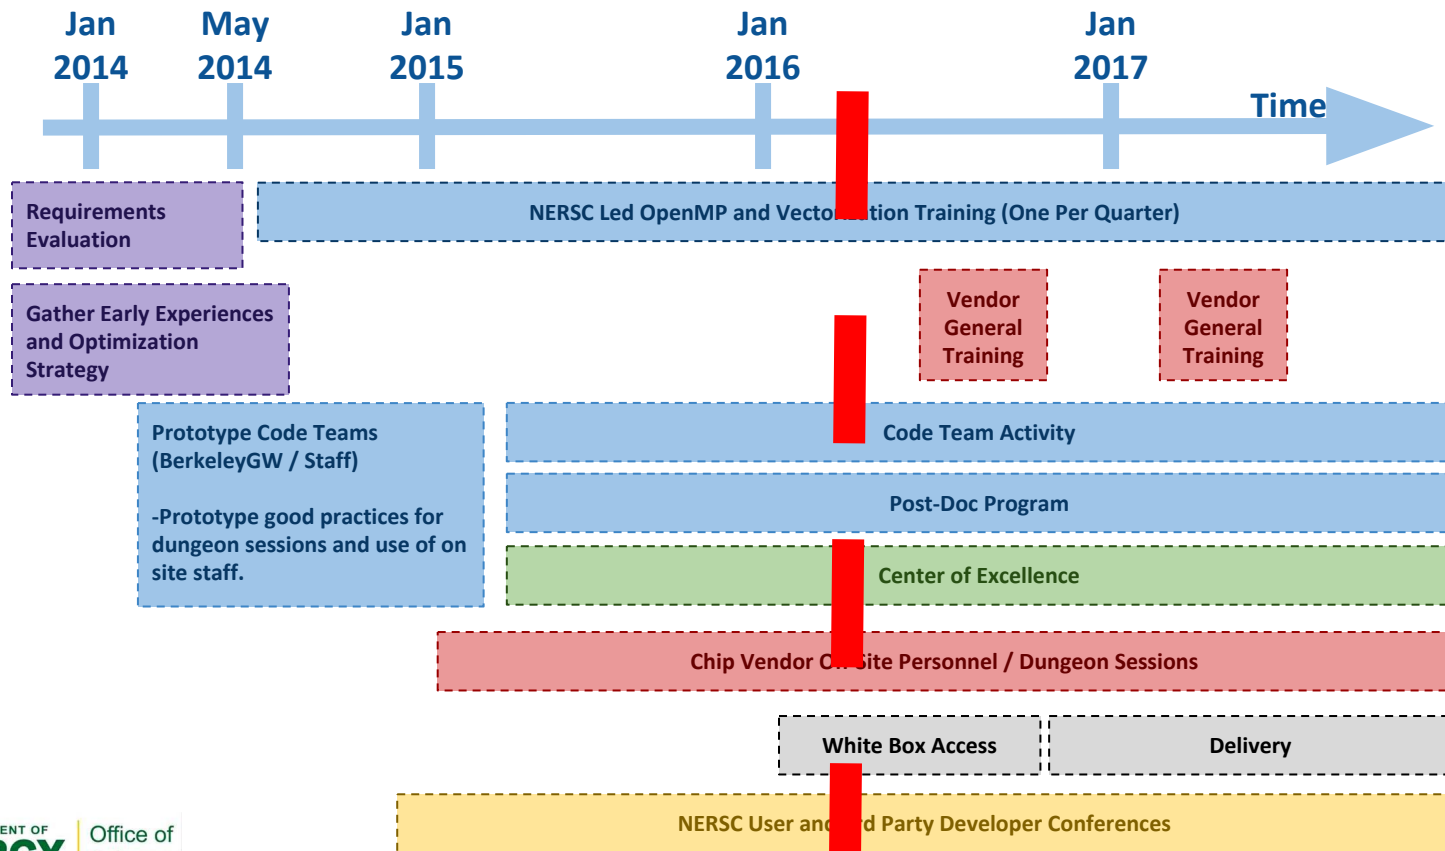
Target Application Team Concept



Timeline

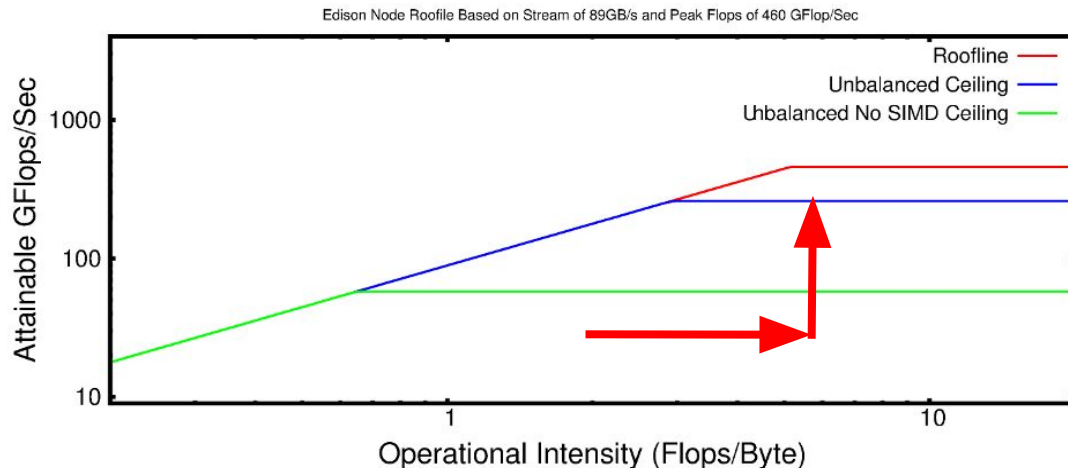


Timeline



Important Optimization Concepts

- MPI+X (Where X is MPI, OpenMP, PThreads, PGAS etc)
- Understanding Memory Bandwidth
- Vectorization



Optimizing Code For Cori is like:

A. **A Staircase ?**

B. **A Labyrinth ?**

C. **A Space Elevator?**



*(More)
Optimized Code*



The Ant Farm!

OpenMP
scales only to 4
Threads

large cache
miss rate

Code shows no
improvements
when turning on
vectorization

50% Walltime
is IO

Communication
dominates beyond
100 nodes



Compute intensive
doesn't vectorize

Memory bandwidth
bound kernel

IO bottlenecks

MPI/OpenMP
Scaling Issue

Can you
use a
library?

Increase
Memory
Locality

Utilize High-Level
IO-Libraries. Consult
with NERSC about
use of Burst Buffer.

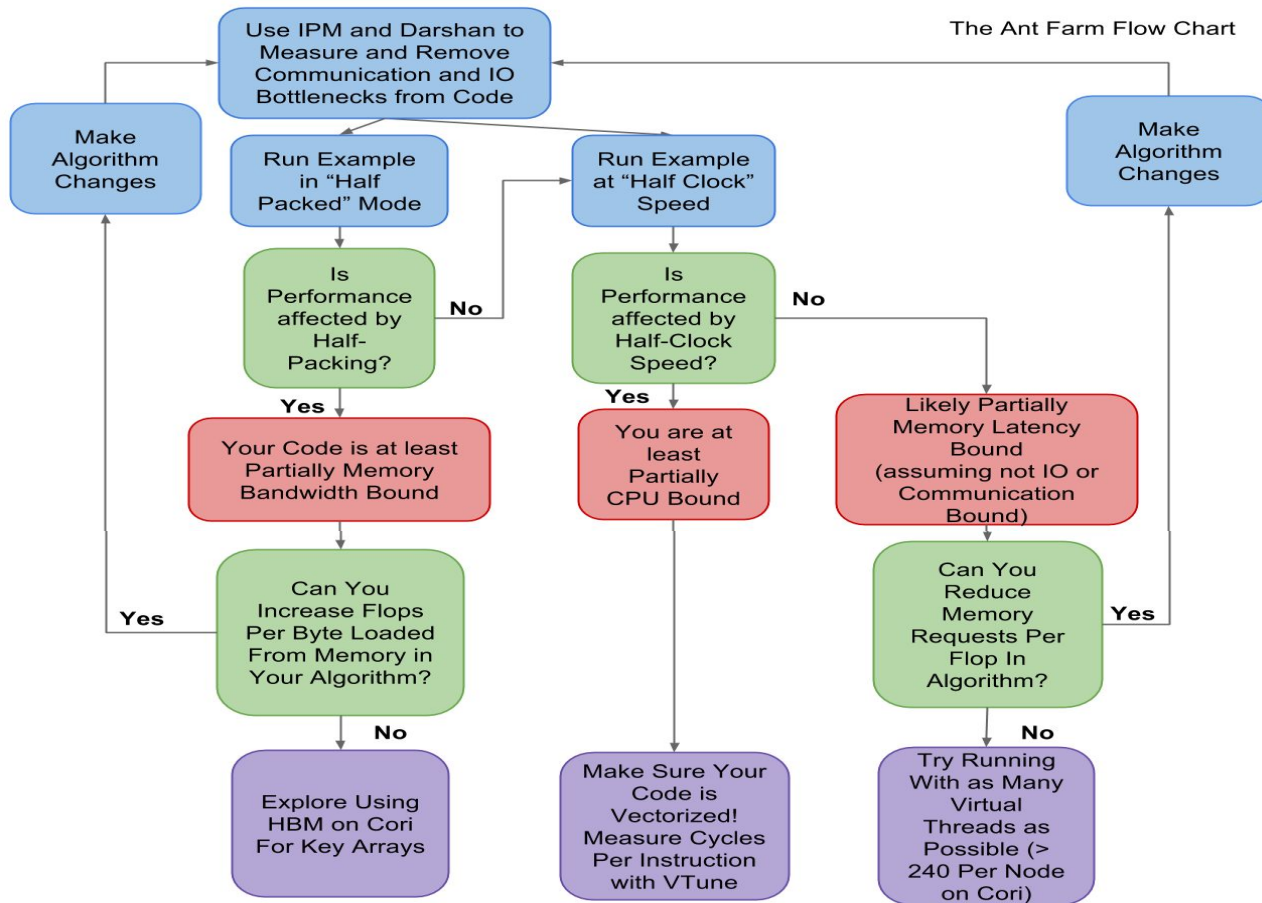
Use Edison to
Test/Add OpenMP
Improve Scalability.
Help from NERSC/Cray
COE Available.

Create micro-kernels or
examples to examine
thread level
performance,
vectorization, cache
use, locality.

Utilize
performant /
portable
libraries

The Dungeon:
Simulate kernels on KNL.
Plan use of on package
memory, vector
instructions.

Dungeon Prep Flow Chart

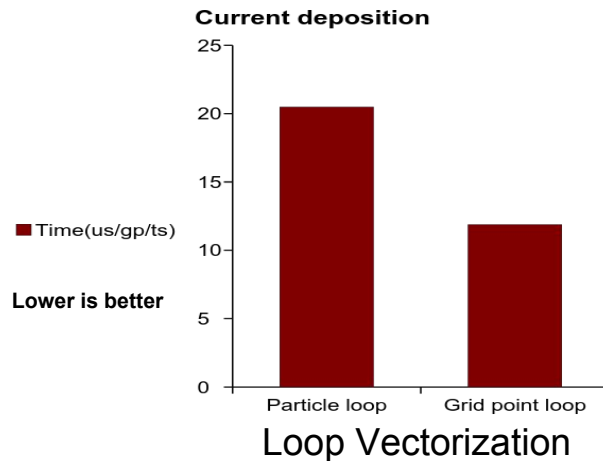
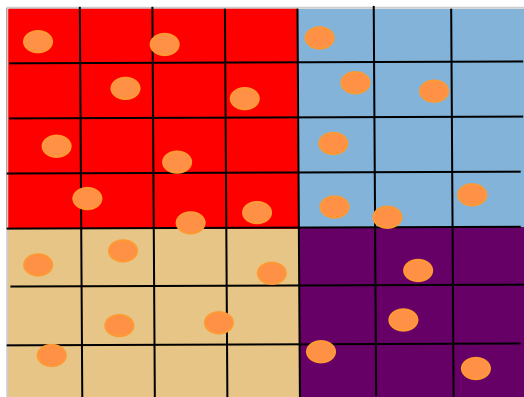


What Has Gone Well

1. Setting requirements for Dungeon Session motivates teams to get started early and improves quality of dungeon session.
2. Engagement with IXPUG and user communities (Exascale Workshops at CRT)
3. Large number of NERSC and Vendor Training (Vectorization, OpenMP, Tools/Compilers)
4. Learned a Lot about Tools and Architecture (VTune, SDE, HBM, Crapat, Reveal etc.)
5. Vendors Keen to Help

Warp Vectorization Improvements at The Dungeon - Directly enabled by tiling work with Cray COE in Pre-dungeon

Tiling
improves
memory
locality



What Our Users Want



- **Performance** - They want to do as much science possible on the largest/most-impactful systems they can
- **Portability** - They run on multiple systems in the DOE and elsewhere. Where possible want fewer branches.
- **Continuity** - They want to know investments they make now won't have to be remade every two years.

Extras

NERSC Staff associated with NESAP



Katie Antypas



Nick Wright



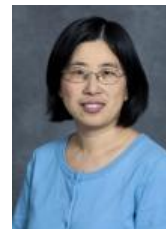
Richard Gerber



Brian Austin



Zhengji Zhao



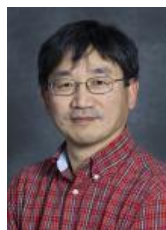
Helen He



Ankit Bhagatwala



Stephen Leak



Woo-Sun Yang



Rebecca Hartman-Baker



Doug Doerfler



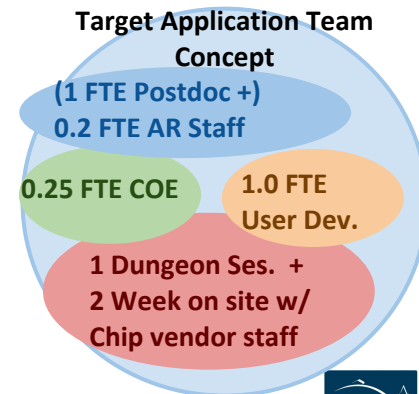
Jack Deslippe



Brandon Cook



Thorsten Kurth



Working With Vendors

NERSC and other centers are uniquely positioned between HPC Vendors and HPC Users and Applications developers.

NESAP provides a power venue for these two groups to interact.

